

CLAIMS:

1 *Sub C1* ① A host coupled to a switched fabric including one or more fabric-attached I/O
2 *St* controllers, comprising:
3 a processor;
4 a host memory coupled to said processor; and
5 a host-fabric adapter coupled to said processor and provided to interface with said
6 switched fabric, which caches selected translation and protection table (TPT) entries from said
7 host memory for a data transaction, and flushes individual cached translation and protection table
8 (TPT) entry in accordance with a translation cacheable flag.

9 2. The host as claimed in claim 1, wherein said host-fabric adapter comprises an
10 internal cache for storing said selected translation and protection table (TPT) entries from said
11 host memory.

12 *Sub C1* 3. The host as claimed in claim 2, wherein each of said selected translation and
13 *Bo* protection table (TPT) entries represents translation of a single page of said host memory.

14 4. The host as claimed in claim 2, wherein said host-fabric adapter is provided to
15 perform virtual to physical address translations and validate access to said host memory using said

1 (c) selected translation and protection table (TPT) entries.

1 5. The host as claimed in claim 2, wherein each of said translation and protection
2 table (TPT) entries comprises:

3 protection attributes which control read and write access to a given memory region of said
4 host memory;

5 said translation cacheable flag which specifies whether said host-fabric adapter may flush a
6 corresponding translation and protection table (TPT) entry stored in said internal cache;

7 a physical page address field which addresses a physical page frame of data entry; and

8 a memory protection tag which specifies whether said host-fabric adapter has permission
9 to access said host memory.

10 cut
11 C1 6. The host as claimed in claim 5, wherein said protection attributes comprise a
12 Memory Write Enable flag which indicates whether said host-fabric adapter can write to page; a
13 RDMA Read Enable flag which indicates whether the page can be source of RDMA Read
14 operation; a RDMA Write Enable flag which indicates whether the page can be target of RDMA
15 Write operation.

1 7. The host as claimed in claim 5, wherein said host-fabric adapter flushes a
2 designated cached translation and protection table (TRT) entry from said internal cache when said

21
1 translation cacheable flag of said designated cached translation and protection table (TPT) entry
2 indicates a first logic state, and maintains said designated cached translation and protection table
3 (TPT) entry in said internal cache when said translation cacheable flag of said designated cached
4 translation and protection table (TPT) entry indicates a second logic state opposite of said first
5 logic state.

1 8. The host as claimed in claim 1, further comprising an operating system including
2 driver software which sets status of said translation cacheable flag per translation and protection
3 table (TPT) entry for enabling said host-fabric adapter to flush individual cached translation and
4 protection table (TPT) entry from said internal cache.

66 1 2 3 4 5 6 7
1 9. A network, comprising:
2 a switched fabric;
3 I/O controllers attached to said switched fabric; and
4 a host comprising an operating system, a host memory, and a host-fabric adapter which
5 caches selected translation and protection table (TPT) entries from said host memory, and which
6 flushes individual cached translation and protection table (TPT) entry in accordance with a
7 translation cacheable flag.

1 10. The network as claimed in claim 9, wherein said host-fabric adapter comprises an

1 internal cache for storing said selected translation and protection table (TPT) entries from said
2 host memory.

1 11. The network as claimed in claim 9, wherein each of said selected translation and
2 protection table (TPT) entries represents translation of a single page of said host memory.

1 12. The network as claimed in claim 9, wherein said host-fabric adapter is provided to
perform virtual to physical address translations and validate access to said host memory using said
selected translation and protection table (TPT) entries.

13. The network as claimed in claim 10, wherein each of said translation and
protection table (TPT) entries comprises:

protection attributes which control read and write access to a given memory region of said
host memory;

said translation cacheable flag which specifies whether said host-fabric adapter may flush a
corresponding translation and protection table (TPT) entry stored in said internal cache;

a physical page address field which addresses a physical page frame of data entry; and

a memory protection tag which specifies whether said host-fabric adapter has permission
to access said host memory.

1 *Sub* 14. The network as claimed in claim 13, wherein said protection attributes comprise a
2 *CI* Memory Write Enable flag which indicates whether said host-fabric adapter can write to page; a
3 RDMA Read Enable flag which indicates whether the page can be source of RDMA Read
4 operation; a RDMA Write Enable flag which indicates whether the page can be target of RDMA
5 Write operation.

1 15. The network as claimed in claim 10, wherein said host-fabric adapter flushes a
designated cached translation and protection table (TPT) entry from said internal cache when said
translation cacheable flag of said designated cached translation and protection table (TPT) entry
indicates a first logic state, and maintains said designated cached translation and protection table
(TPT) entry in said internal cache for future re-use when said translation cacheable flag of said
designated cached translation and protection table (TPT) entry indicates a second logic state
opposite of said first logic state.

1 *Sub* 16. An apparatus which stores translation and protection table (TPT) entries for virtual
2 *AK* to physical address translations, and which flushes individual translation and protection table
3 (TPT) entry stored in accordance with a corresponding translation cacheable flag.

1 17. The apparatus as claimed in claim 16, further comprising an internal cache for
2 storing said translation and protection table (TPT) entries.

1 18. The apparatus as claimed in claim 16, wherein each of said translation and
2 protection table (TPT) entries represents translation of a single page of a host memory.

1 19. The apparatus as claimed in claim 17, wherein each of said translation and
2 protection table (TPT) entries comprises:

3 protection attributes which control read and write access to a given memory region of a
4 host memory;

5 said translation cacheable flag which specifies whether said apparatus may flush a
6 corresponding translation and protection table (TPT) entry stored in said internal cache;

7 a physical page address field which addresses a physical page frame of data entry; and

8 a memory protection tag which specifies whether said apparatus has permission to access
9 said host memory.

1 20. The apparatus as claimed in claim 19, wherein said protection attributes comprise a
2 Memory Write Enable flag which indicates whether said apparatus can write to page; a RDMA
3 Read Enable flag which indicates whether the page can be source of RDMA Read operation; a
4 RDMA Write Enable flag which indicates whether the page can be target of RDMA Write
5 operation.

add a2